The Role of Gold in Silicon Thyristors

David C. Northrop
University of Manchester Institute of Science and Technology, U.K.

Silicon thyristors are widely applied in controlling the power input to devices ranging in scale from domestic light dimmers to the megawatt electric motors in steel rolling mills. This article outlines the principles of operation of thyristors and highlights the role of gold, first in achieving their rapid transition from the conducting to the non-conducting state, and secondly in the simultaneous creation of n-type layers on the silicon and of electrical contacts to these layers.

During the early development of solid state electronics in the 1950's, it was widely held that vacuum tubes would never be supplanted by semiconductor devices for either high power or high frequency applications. Both these predictions have proved to be spectacularly wrong. In the field of microwave engineering, devices such as Gunn and impact avalanche and transit time (IMPATT) oscillators have produced power at higher frequencies than is possible from vacuum tubes, and at levels sufficiently high for radar systems. In many ways the devices known as thyristors, which are used in a wide variety of power control applications, are even more impressive, for they handle very much larger amounts of power in a single component than any thermionic valve ever handled, and with much greater efficiency and reliability. Add to this list the fact that they are much cheaper than their vacuum tube equivalents, and it is not surprising that they have displaced the latter and are now used for power control in an extended range of situations.

This last point is very important and can easily be illustrated by reference to domestic appliances. Many homes in the developed countries contain thyristors which control the power input to cookers and washing machines, electric light dimmers and television sets. Most of these applications require compact devices which do not themselves dissipate appreciable heat and do not have an extended 'warming up' time after switching on. For these reasons, designs based on vacuum tubes would not be suitable, even if the cost and reliability of such tubes were acceptable. Thus, the demand for thyristors is constantly increasing, and this in turn helps to reduce costs through mass production.

It should also be noted that the physical size of thyristors is very large compared with that of most other semiconductor devices. Typically, one of the many transistors in a modern integrated circuit is a few micrometres in size, and the whole integrated circuit, consisting of literally thousands of transistors and their interconnection patterns, is made on a silicon chip only a few square millimetres in area. In contrast, a single thyristor rated at 4 000 V and 3 000 A, might require a whole slice of silicon 50 mm in diameter. However, although it is right to make this comparision in order to emphasize the enormous difference of scale between thyristors and transistors, it is equally important to compare the former with the old mercury vapour rectifiers and thyratrons of earlier power control technology. These were often more than a metre high, and handled less power than a single thyristor.

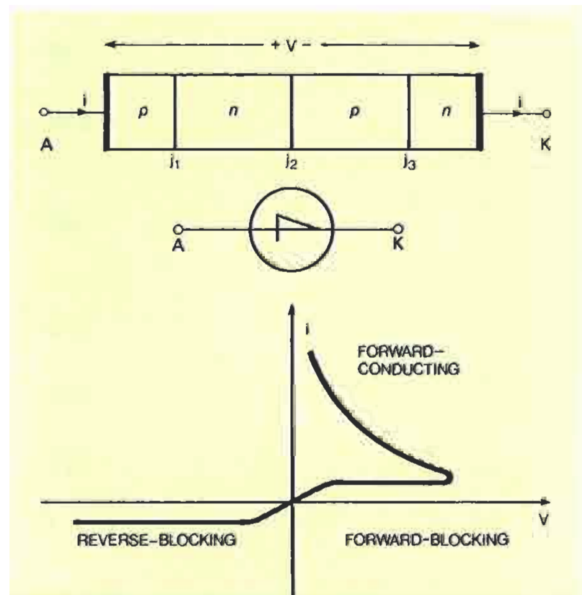

Fig. 1 Schematic representation of the four-layer *p-n-p-n* diode, its conventional circuit symbol and its current-voltage characteristics. A and K denote anode and cathode respectively, and j_1, j_2 and j_3 the junctions between layers of the semiconductor materials

In describing the operating principles of thyristors it is helpful to follow their historical development, beginning with the original *p-n-p-n* diode (1), then progressing to the silicon controlled rectifier (SCR) in which a third, control, electrode is added and then to the more sophisticated diacs and triacs.

The *p-n-p-n* Diode

The four-layer *p-n-p-n* structure shown in Figure 1 has only two terminals, labelled anode A and cathode K, but there are three rectifying junctions: j_1, j_2 and j_3.

The current i and voltage V are shown corresponding to the forward direction of conduction, in which the junctions j_1 and j_3 are forward-biased (that is, they present an insignificant resistance to current flow in the direction required by the applied voltage). The centre junction j_2 is reverse-biased and presents the only barrier to current flow, so that it accounts for the total voltage drop across the diode. If V is increased, initially the applied voltage is absorbed without appreciable increase in current, but eventually the whole device switches to a new condition, in which a much higher current can flow at a much lower applied potential. When this happens, the current is controlled by an external circuit resistor across which most of the available voltage now appears. It is this new conducting state, in which all three junctions behave as forward-biased, that we must understand if we are to appreciate how all the other members of the thyristor family operate. Before examining this behaviour in greater detail, we should note that in the reverse direction of applied voltage, where junctions j_1 and j_3 are reverse-biased, the flow of current is effectively blocked and as the voltage is increased, nothing particularly interesting happens.

The switch from the forward-blocking to the forward-conducting state can be understood qualitatively in terms of the two-transistor model, illustrated in Figure 2. This shows the device of Figure 1 broken down into two interconnected 'bipolar' transistors, transistor 1 being of the *p-n-p*-type and transistor 2 of the *n-p-n*-type. The emitters E, bases B, and collectors C, are labelled, as are the current flows between the two separate devices. We now have to make a small detour from our main path in order to understand qualitatively how a bipolar transistor operates (2).

The Bipolar Transistor

The bipolar transistor is a three-terminal, two-junction device designed to have one of its two junctions reverse-biased irrespective of the polarity of the applied voltage. Each individual junction, or diode, is asymmetric electically, as illustrated in Figure 3(a) where the directions of current and voltage are shown positive for forward conduction. In *p*-type silicon the

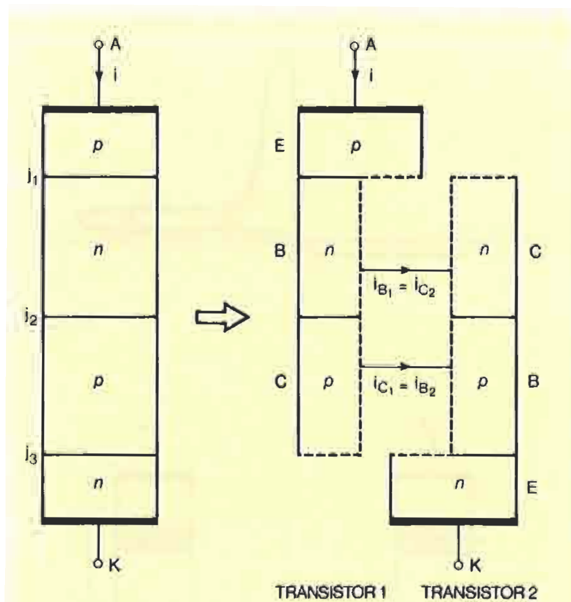

Fig. 2 The two-transistor model of a *p-n-p-n* diode. The collector current, i_{C_1}, of the *p-n-p* transistor 1 becomes the base current, i_{B_2}, of the *n-p-n* transistor 2, and similarly the base current of transistor 1, i_{B_1}, is the collector current of transistor 2, i_{C_2}. The centre junction j_2 is common to both transistors

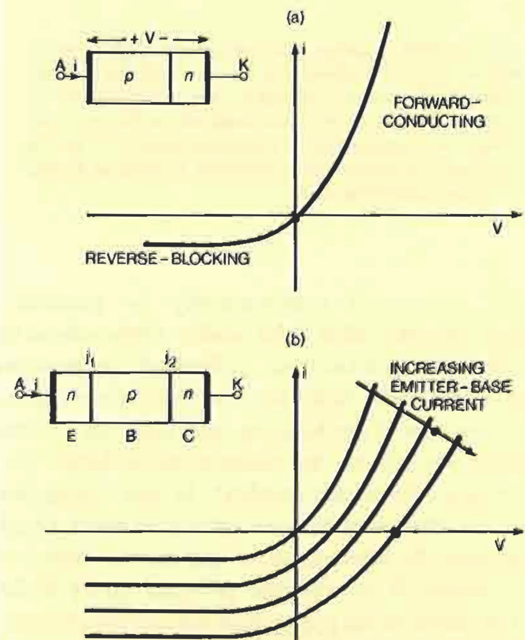

Fig. 3 Schematic description of transistor action. The current-voltage characteristics of (a) a single *p-n* junction, and (b) a *p-n* junction modified by 'injection' from a second *p-n* junction in close proximity to it

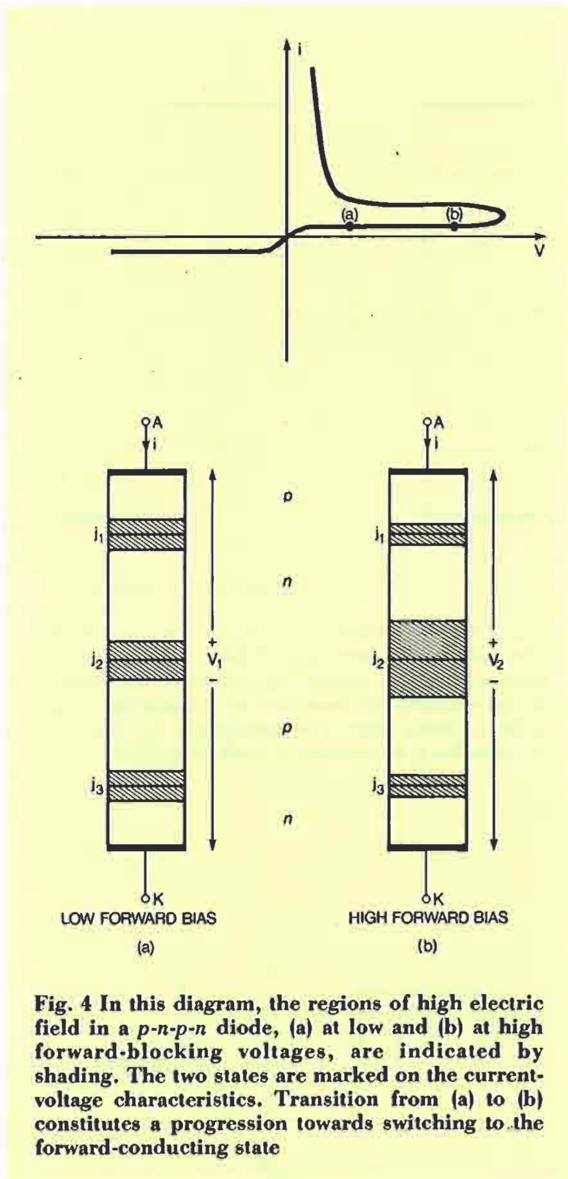

Fig. 4 In this diagram, the regions of high electric field in a *p-n-p-n* diode, (a) at low and (b) at high forward-blocking voltages, are indicated by shading. The two states are marked on the current-voltage characteristics. Transition from (a) to (b) constitutes a progression towards switching to the forward-conducting state

current is carried predominantly by positively-charged particles called holes, and in *n*-type silicon by negatively-charged electrons. A forward bias provides forces which cause holes to cross from the *p*-type to the *n*-type side of the junction, and electrons to cross the other way (giving two components of electric current which *add* to one another). In this conduction mode, the charge-carrier concentrations remain high throughout the semiconductor and current therefore flows readily. A reverse bias provides forces which repel the holes in the *p*-type material and the electrons in the *n*-type material away from the junction, so producing a region of low charge-carrier concentration through which current flows only with great difficulty. The junction is able to withstand considerable reverse voltages without an appreciable increase in current. Such current that does flow is composed of the very small number of free electrons in the *p*-type material and the similarly small number of free holes in the *n*-type material (the so-called 'minority carriers'). These two groups of charge-carriers are assisted in crossing the junction by the reverse bias field, and the current which flows is much less dependent on the applied voltage than on the supply rate of minority carriers. In fact, the current is determined by the rates at which these are generated by thermal ionization in a small volume of the semiconductor in the vicinity of the *p-n* junction.

A transistor behaves differently from two separate diodes because of the close proximity of the two junctions. Their interaction is illustrated in the current-voltage characteristics shown in Figure 3(b). When the base-collector junction j_2 is reverse-biased to the applied voltage, it conducts a small current, as explained above, but at the same time the emitter-base junction, j_1, is forward-biased. The current through j_1 consists in part of electrons injected into the transistor base, and if any of these reach j_2 they add to the available supply of minority carriers in the *p*-type material, thereby increasing the reverse current through that latter junction. This effect of one junction's current modulating the current in the other junction is referred to as 'transistor action'. This behaviour can be expressed quantitatively by the common base current gain, α, which measures the fraction of the emitter-base junction current, i_e, reaching the base-collector junction as i_c. Thus:

$$i_c = \alpha\, i_e + I_{co} \qquad (1)$$

where I_{co} is the unmodified reverse current of the base-collector junction.

The Thyristor in Forward Conduction

All the necesary elements have now been assembled to allow the further discussion of the forward-conducting state of the *p-n-p-n* diode via the two-transistor model shown in Figure 2. If transistors 1 and 2, into which the device is divided, have current gains α_1 and α_2 respectively, a simple algebraic manipulation based on equation (*1*) and the current equalities given in Figure 2 shows that:

$$i = \frac{I_{co_1} + I_{co_2}}{1 - (\alpha_1 + \alpha_2)} \qquad (2)$$

If $(\alpha_1 + \alpha_2)$ is much less than unity, then the current is small and roughly equal to the currents through the two reverse-biased junctions, but if $(\alpha_1 + \alpha_2)$ can be made to approach unity, then the current becomes very large indeed. In a single transistor, it is possible for α to be almost equal to unity (its maximum value).

This is achieved by positioning the two junctions very close to one another — in practice, a few micrometres apart. This is not possible for all three junctions of a *p-n-p-n* diode, nor is it particularly desirable. Thus, the process of switching from forward-blocking to forward conduction requires that the component transistors have values of α that are low at low voltage, but which rise with increasing voltage. There are several factors which contribute to the increase in $(\alpha_1 + \alpha_2)$, but the most important is the widening of the region of high electric field around the base-collector junction of transistor 1. The high field results from the removal of holes and electrons by the applied voltage, and is necessary to maintain the same (very low) current density throughout the device. This high field region is, in effect, an extension of the base-collector junction and, as shown in Figure 4, its widening leads to a reduction in the separation between junctions j_1 and j_2.

Once switching has occurred, this state of affairs changes, especially near the central junction. Figure 5 shows the charge and field distributions before and after switching has taken place. Before switching, the charge distribution around j_2, shown in Figure 5(a), consists almost entirely of fixed, ionized impurities and the resulting large electric field (see Figure 5(b)) can be deduced from the charge distribution by using Gauss's Law. The mobile charges (electrons and holes) are not shown. They are very few and are distributed in such a way that the current produced by the field acting on them is the same everywhere in the device. In the forward-conducting state, the influence of additional mobile charges becomes important near the central junction j_2. The forward-biased junction j_1 pumps holes into the *n*-type region adjacent to j_2, and at the transition to the forward-conducting state many of these holes reach the high field region near j_2. This is shown in Figure 5(c), where it is important to note that an accompanying equal number of electrons is present in the vicinity of the central junction. This is relevant to the operation of the transistor, as well as the thyristor, and stems from the comparative ease with which free charges of both signs can move to maintain charge neutrality (2). These additional mobile charges clearly require a low electrical field to cross the junction j_2 for a given current, and so the voltage across this junction falls in a way which is consistent with maintaining constant current and with the new charge distribution which takes into account the mobile charges and the fixed charges shown in Figure 5(c) and 5(a) respectively. These fields and voltage reductions, (see Figure 5(d)) have a regenerative effect, because if the circuit voltage is held constant, a greater potential difference occurs across j_1 and j_3, increasing the number of mobile charges near j_2. It is a surprising, but nevertheless a true, observation that in the forward-conducting state so much charge has been injected into j_2 from both sides, that it is in forward bias, that is, the potential difference across it opposes the external circuit voltage.

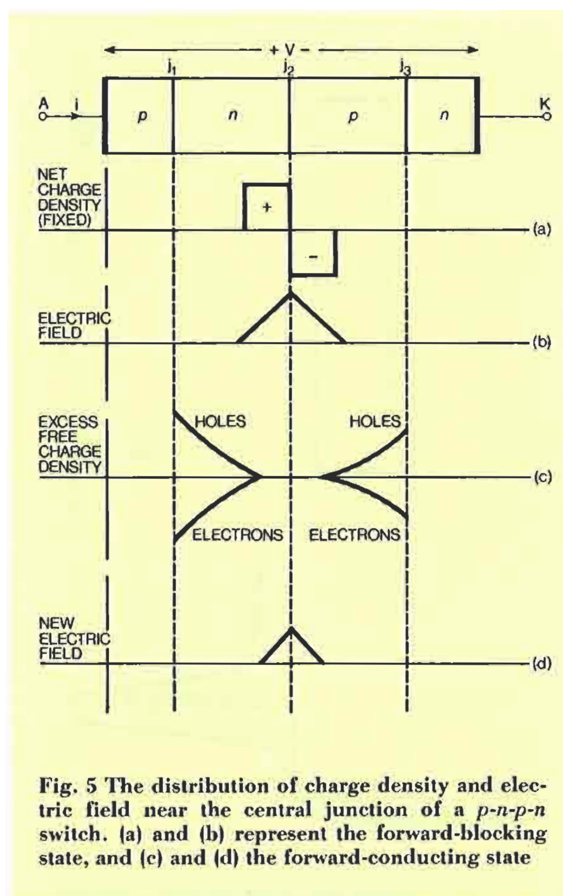

Fig. 5 The distribution of charge density and electric field near the central junction of a *p-n-p-n* switch. (a) and (b) represent the forward-blocking state, and (c) and (d) the forward-conducting state

This rather lengthy discussion of the *p-n-p-n* switch has established a foundation from which more useful devices can be described, and the important role of gold atoms which are introduced into the device structure can be appreciated. Briefly, we shall consider the ways in which a third connection to the device is used in the SCR and the triac, to produce controlled triggering of the forward-conducting state, and then, having seen how the devices are switched on, we must consider how they are to be switched off. This latter problem is a serious one. As already discussed, the blocking and conducting states are characterized by markedly differing field and charge distributions (the two distributions are linked in a complicated way by Gauss's Law and the steady-state current continuity equation). In order to turn off the forward conduction it is necessary to eliminate the extra charge-carriers that have accumulated near the central junction, and this is made possible by the presence of minute amounts of gold that have been diffused into the semiconductor lattice.

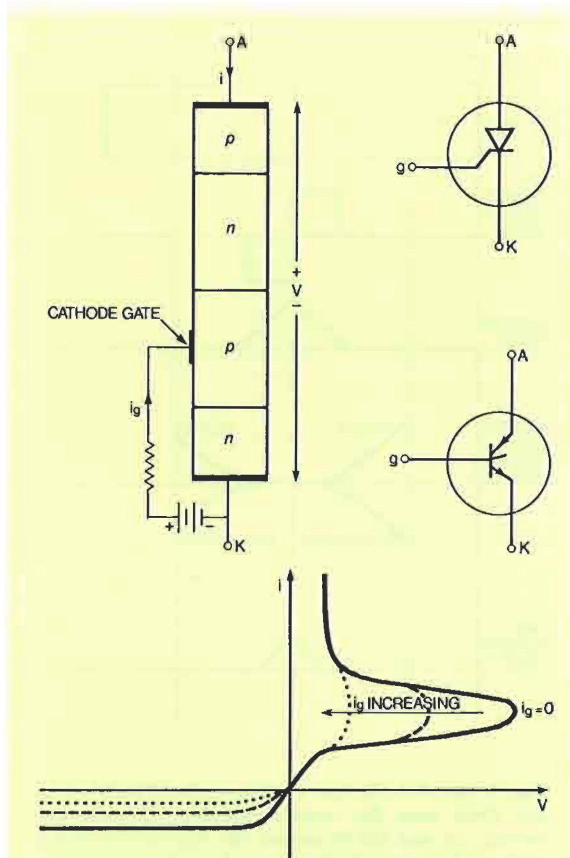

Fig. 6 The cathode gate semiconductor controlled switch, its conventional circuit symbols and the current-voltage characteristics for increasing values of the grid current, i_g. Changes in the very low grid current control the anode-to-cathode voltage at which switching from the blocking to the conducting state occurs

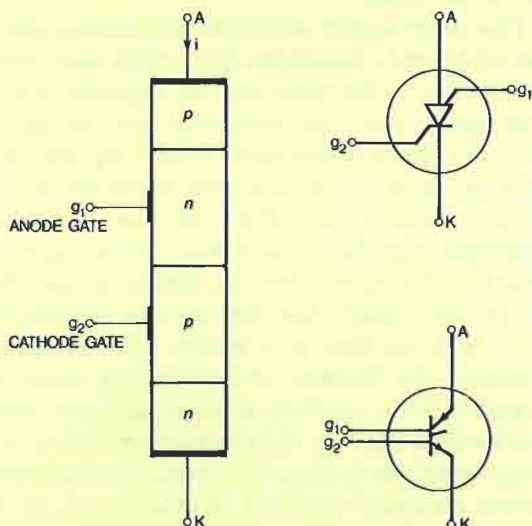

Fig. 7 Schematic representation and circuit symbols for the semiconductor controlled switch (SCS). This device is fitted with both cathode and anode grid electrodes and may be triggered via either of them

The SCR and the SCS

The usefulness of the thyristor is greatly increased by the introduction of a third electrode, the so-called grid or gate, which is connected to the *p-n-p-n* structure to form an SCR, as shown schematically in Figure 6. The transition from the forward-blocking to the forward-conducting state is assisted by the positive current flow through the gate electrode (1). This has exactly the same effect on the charge density near j_2 as that achieved by increasing the total anode-to-cathode voltage and causes switching between states at much lower applied voltages. Clearly, the SCR is far more useful than the simpler *p-n-p-n* diode, since the conducting state, capable of passing many amperes can be controlled by a gate current, i_g, of a few milliamperes passing for a few microseconds. However, once the forward conduction has been established, the gate electrode loses control and cannot be used to switch the device back to the forward-blocking state.

Switching from blocking to conducting can also be achieved by a negative current pulse to a gate electrode near the anode, and some devices are made with both anode and cathode gates (Figure 7). These are known as semiconductor controlled switches (SCS) and have greater flexibility in use than SCR's, since they can be triggered in either of two ways, but their power-handling capacity is not increased by the presence of two grid electrodes.

Bilateral Devices, the Diac and Triac

Many applications call for devices which are capable of switching readily in either direction. This can be achieved straightforwardly and effectively by diacs which are, in essence, two *p-n-p-n* diodes in parallel. Figure 8 shows such a device schematically, together with its (symmetrical) current *vs.* voltage characteristics. One diode is formed by the layers p_1, n_1, p_2 and n_2 and the other by p_2, n_1, p_1 and n_3. Junctions j_4 and j_3 are short-circuited by the anode and cathode contacts. Although geometrically more complicated than *p-n-p-n* diodes, diacs switch from the blocking to the conducting state by the same physical mechanism.

The next stage of complexity is to introduce one or two gate electrodes to trigger the conducting state in either, or both, directions of applied voltage. Again, the mechanism is familiar; it is indentical to that of the SCR. These bilateral triode switches are sometimes called triacs.

Switching Off — The Role of Gold

As discussed above, the charge distribution in the thyristor in the blocking or 'off' state is quite different from that in the conducting or 'on' state. In the latter,

there exists considerably more free charge, consisting of approximately equal numbers of negative electrons and positive holes, near the central junction j_2 of the *p-n-p-n* diode, and in order to switch off the device, all this charge must be removed. Since the excess charge is mobile, it can be removed by the application of appropriately directed electric fields, for example by reversing the gate or anode potentials. This method has its usefulness, but the gate electrode is not particularly suited to this purpose, because current has to be drawn laterally through the device structure over comparatively large distances. The technique of anode circuit commutation (that is, reversing the field directions in the device) is much more effective, because the charge extraction involves much shorter distances. There are, however, many applications where the method is not possible, simply because external circuit conditions do not allow the reversal of the anode voltage, or where the aditional circuit complexity required for this operation is not permissible.

The solution of the problem lies in an entirely different approach. Rather than to remove them from the device, pairs of electrons and holes may be 'encouraged' to recombine with one another in an annihilation process that reduces charge concentration *in situ* (3,4). Since the excess electron and hole concentrations are nearly equal, this process can be effective in restoring almost completely the charge distribution of the forward-blocking state. In order to understand this, we must digress slightly to consider the formation and removal of electrons and holes in a semiconductor.

Figure 9(a) shows in its simplest form the conventional energy band diagram of a semiconductor. The valence band, at the energy level E_v, contains N_v 'states' which can be occupied by a maximum number of N_v electrons. At very low temperatures, all these states are filled and the corresponding complete set of N_v electrons does not conduct electricity. The reason for this interesting and important fact is that the complete set of momenta for the electrons adds vectorially to zero, that is, there is no net motion of charge. In a semiconductor, the conduction band at the energy level $E_c = E_v + E_g$ is empty of electrons at very low temperatures and, therefore, it too makes no contribution to the electrical conductivity. When an electron is excited from the valence band to the conduction band, as shown in Figure 9(a) by transition 1, electrical conduction can take place. The electron can now move freely through the material in response to applied electric fields, and also, there can occur a rearrangement of the $N_v - 1$ electrons remaining in the valence band amongst the N_v available states, both of which allow current flow in response to externally applied fields. The hole is a particle of positive charge and positive mass, the properties of which are

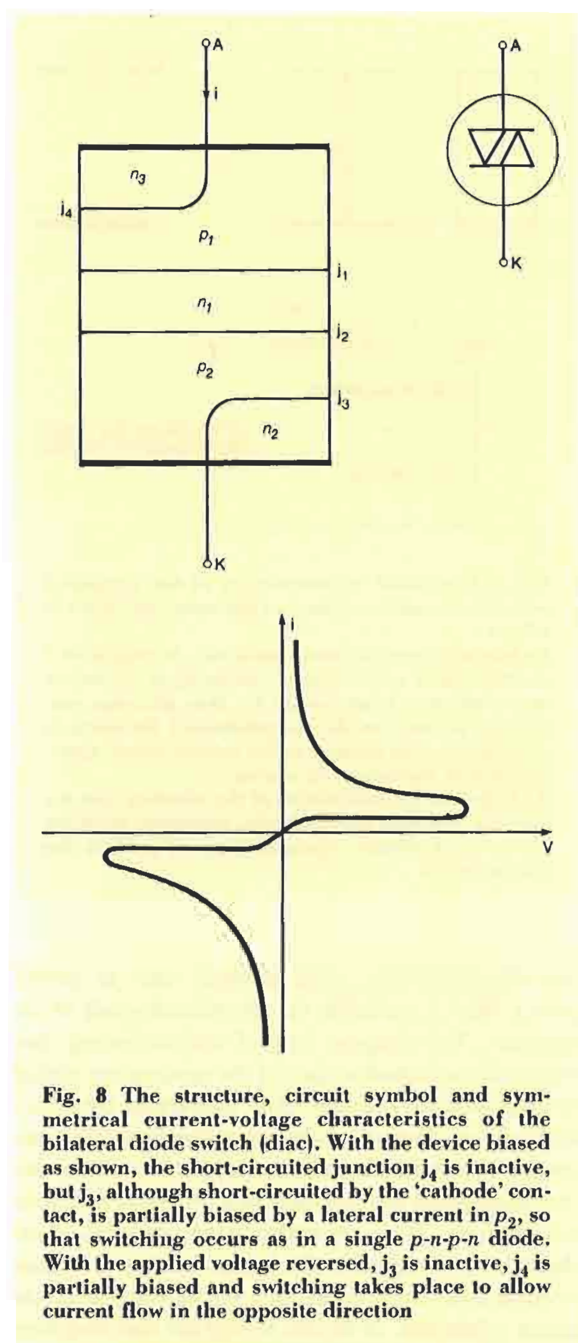

Fig. 8 The structure, circuit symbol and symmetrical current-voltage characteristics of the bilateral diode switch (diac). With the device biased as shown, the short-circuited junction j_4 is inactive, but j_3, although short-circuited by the 'cathode' contact, is partially biased by a lateral current in p_2, so that switching occurs as in a single *p-n-p-n* diode. With the applied voltage reversed, j_3 is inactive, j_4 is partially biased and switching takes place to allow current flow in the opposite direction

so defined that its motion in applied electric fields results in exactly the same current flow as that produced by the $N_v - 1$ electrons. For convenience, transition 1 in Figure 9(a) is referred to as the creation of an electron-hole pair and transition 2 as recombination. The latter restores the material to a non-conducting state by the simultaneous annihilation of an electron and a hole. It is precisely this process which has to be encouraged to switch off the thyristor.

In silicon, the direct recombination process depicted in Figure 9(a) is not very efficient, and the

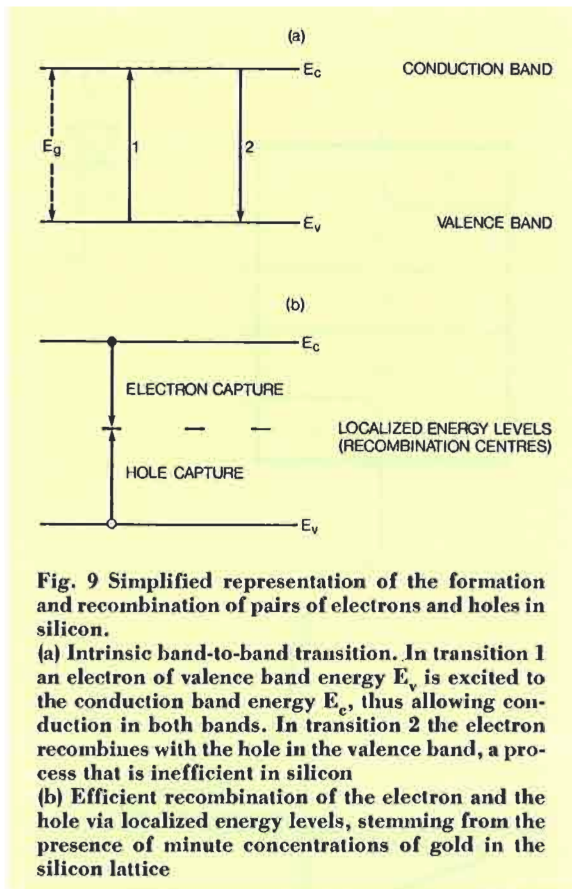

Fig. 9 Simplified representation of the formation and recombination of pairs of electrons and holes in silicon.
(a) Intrinsic band-to-band transition. In transition 1 an electron of valence band energy E_v is excited to the conduction band energy E_c, thus allowing conduction in both bands. In transition 2 the electron recombines with the hole in the valence band, a process that is inefficient in silicon
(b) Efficient recombination of the electron and the hole via localized energy levels, stemming from the presence of minute concentrations of gold in the silicon lattice

free charge-carriers, once created, tend to persist longer than is desirable for the switching off of the thyristor. The simplest way of understanding this, derives from considerations of the momentum carried by each electron and hole. Any interaction between a hole and an electron, including recombination, must conserve momentum, and in silicon it is unlikely that the recombined electron-hole pair will have zero total momentum. The recombination process must therefore be coupled to the lattice, so that the excess electron-hole momentum may be transferred to the lattice vibrations. If we can strengthen this coupling, then recombination occurs more readily and the time before the excess electrons and holes disappear is reduced.

The best method of achieving this is shown in Figure 9(b), where additional energy levels are introduced about half way between E_v and E_c. Such levels can be produced by introducing atoms of carefully chosen impurities into the silicon crystal. These energy levels differ from those in the valence and conduction bands by being localized in space, so that electrons in them are confined within the volume occupied by a few atoms (the so called recombination centres). In contrast, the energy levels of the conduction and valence bands are non-localized, and electrons occupying them can move freely through the whole of the crystal. One important consequence of this localization is a strong coupling between an electron 'captured' at a recombination centre and the lattice. Another equally important aspect is that the separation of the recombination process into the two stages shown in Figure 9(b) makes the probability of recombination greater. Speaking very loosely, it may be said that by localizing an electron we make it easier for the hole to find it. (Recombination may also take place in the reverse sequence: hole capture at a recombination centre, followed by electron capture).

In the selection of the impurity atoms which are to act as recombination centres, gold is a fairly obvious choice because of an advantageous concurrence of properties. First, there is the fact that gold atoms do, in practice, act as extremely effective recombination centres in silicon. Secondly, and more surprisingly, gold atoms are very small and can therefore be introduced into the silicon lattice by solid state diffusion in the relatively low temperature range of 750 to 850°C (5), well below the melting point of silicon (1 420°C). Although gold atoms are very heavy, their small physical size is not anomalous. The mass of an atom is contained almost entirely in its nucleus, whereas its size is determined by the configuration of its electrons. Their compact size allows gold atoms to diffuse interstitially in silicon, that is, to pass between the silicon atoms without displacing them appreciably from their equilibruim positions. It is this ease of movement which makes the introduction of gold possible at relatively low temperatures. Thirdly, the solubility of gold in silicon is fairly low and temperature-dependent (6). This is important, because it is much easier to control the concentration of gold in silicon by saturation than by any other

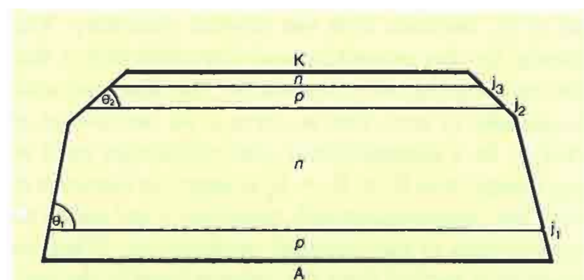

Fig. 10 Sectional view of a typical modern thyristor. The original *n*-type silicon is interdiffused with boron to form the upper and lower *p*-type layers. By alloying the upper of these silicon surfaces with a gold-antimony alloy, the top *n*-type layer is generated, together with an outermost gold-rich contact film. The bevelling of the edges of the device assists high voltage operation. The angles θ_1 and θ_2 are typically 20° and 2°. They appear larger here, because the vertical scale is greatly expanded

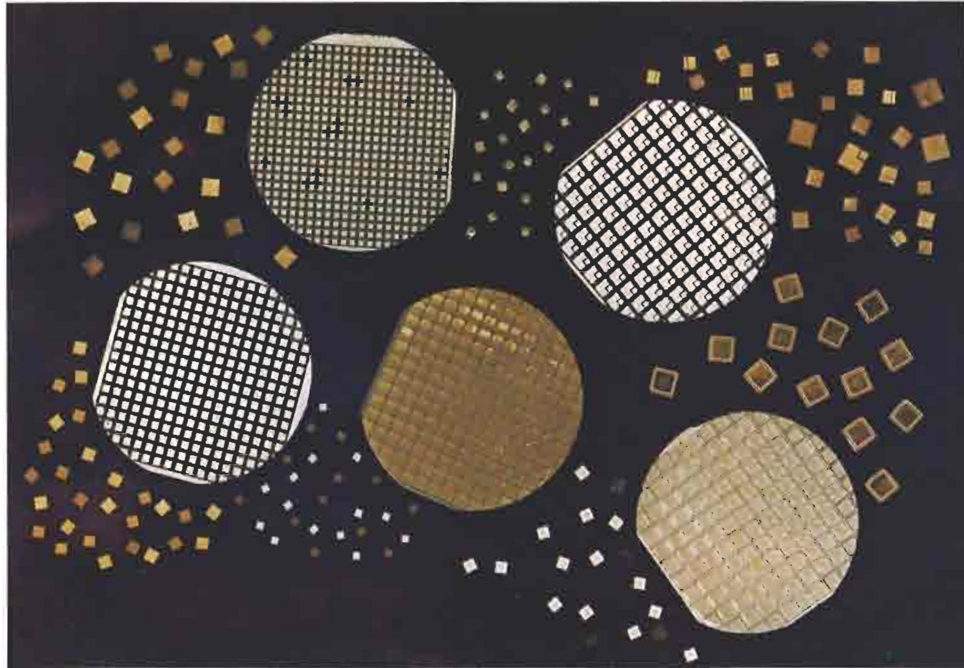

Fig. 11 A variety of processed silicon wafers, 50 mm in diameter, each consisting of a large number of thyristors, together with individual items which have been separated from similar wafers. The thyristors vary in size, according to their current rating. On close inspection the top surface of some devices reveals two separate areas. The larger of these is the cathode which has been formed by the alloying of gold-antimony to the *p*-type layer, previously prepared by diffusion of boron. The smaller area is that of the cathode gate in which gold-gallium has been alloyed with the *p*-type silicon to form an ohmic contact. The gate is smaller than the cathode because it carries much lighter currents

Photograph by courtesy of Mullard Limited

means. Thus, by a careful choice of the diffusion temperature, a uniform concentration of gold can be achieved at a level that is effective for recombination purposes, but which causes no serious side-effect, such as an increase in the electrical resistivity of the silicon.

Structure and Manufacture of Thyristors

The diagrams of thyristors shown so far have been schematic representations intended to illustrate operating principles. Figure 10 shows a section through a modern silicon thyristor.

The starting material for such a device is usually a slice of *n*-type silicon, the electrical properties of which are as uniform as possible over its entire volume. Its resistivity must be chosen to be compatible with the voltage rating of the thyristor which is to be manufactured. The higher the reverse-blocking voltage required, the higher the resistivity, and hence the higher the purity required in the silicon. It can therefore be seen that the demand for higher voltage operation puts severe pressure on the manufacturer of silicon. (Silicon is already produced industrially at a much higher purity than any other material). The area of the device is determined mainly by its current rating. For example, a mean current of 1 000 A would require the whole of a slice 50 mm in diameter, but for lower current ratings the slice could be cut up after processing to give a number of devices. It probably goes without saying that high voltage and high current ratings are very difficult to obtain together.

The first step in the manufacture of a thyristor consists of the preparation of a 0.25 to 0.4 mm thick slice of silicon with clean, optically polished surfaces. Then follows the introduction of boron by solid state diffusion at temperatures of 1 150 to 1 200°C for up to 48 hours, to form the two *p*-type layers on either side of the original *n*-type layer. The penetration of boron increases with diffusion temperature and time, and is allowed to proceed until the junctions j_1 and j_2 are positioned 0.02 to 0.05 mm from the faces of the slice. The precise junction depths must be chosen in conjunction with the original slice thickness to give the required separation between j_1 and j_2. The charge depletion layers around these junctions are required to spread to achieve close proximity to one another, but they may not overlap. The remaining *n*-layer is then produced by alloying to the silicon a layer of gold containing about 1 per cent antimony. It is another convenient fact that gold and silicon form a eutectic alloy having a melting point of only 377°C. When this alloy solidifies, the first layer to form at j_3

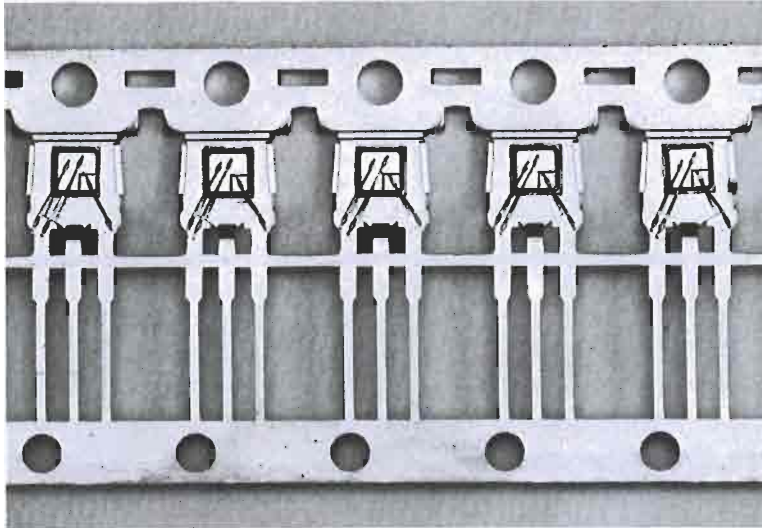

Individual silicon chips have been bonded via the gold-silicon eutectic to a metal frame, thus forming the anode contact. Two gold wires connect each anode to one lead of the frame and a single wire links the cathode gate to another lead. The wires are attached to the frame and to the gold-plated surface of the thyristor by thermocompression bonding

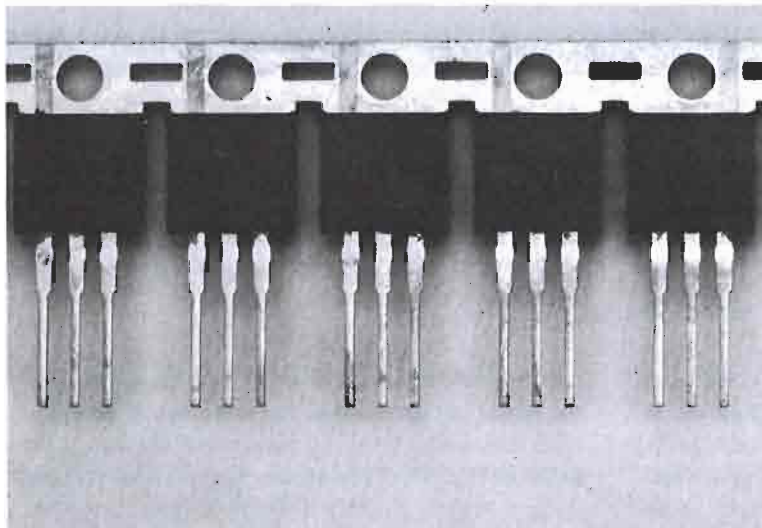

The finished devices have been encapsulated in rectangular blocks of resin and superfluous portions of the frame have been removed. The last operation will consist of cutting the strip into individual components

Fig. 12 Medium power SCR's before (top) and after (bottom) encapsulation. The devices are approximately 3 mm square

Photograph by courtesy of Mullard Limited

is n-type silicon heavily doped with the donor impurity antimony, and the final layer is a highly conductive gold-silicon alloy to which electrical contact can readily be made. The same alloying process may also be used to bond the base of the device to its metal stub; alternatively a gold alloy containing one per cent gallium may be employed. These alloying procedures very conveniently occur at intermediate temperatures — higher than the device operating temperature, so that the structure will always remain solid, but much lower than the diffusion temperatures, so that existing impurity distributions are not disturbed.

It is clear that the introduction of the gold which provides the recombination centres throughout the device must take place before the alloying step. It is, in fact, usually carried out before the boron diffusion step that produces the p-type layers.

The function of gold in a thyristor is thus threefold: it provides the recombination centres required at the central junction for the rapid switching from the conducting to the non-conducting state, it assists in forming the upper n-type layer and it provides on top of the device a layer amenable to bonding.

The bevelling of the edges of the thyristor is important in securing high operating voltages. Because of the difficulty of maintaining complete cleanliness during all stages of manufacture, the surfaces of the silicon may become contaminated in a manner which increases the electrical field in those regions where the high field of any reverse-biased junction intercepts the free surface of the silicon. The bevels are intended to counter this by reducing the surface field.

Figures 11 and 12 are photographs of a range of medium-sized thyristors. Figure 11 shows a number

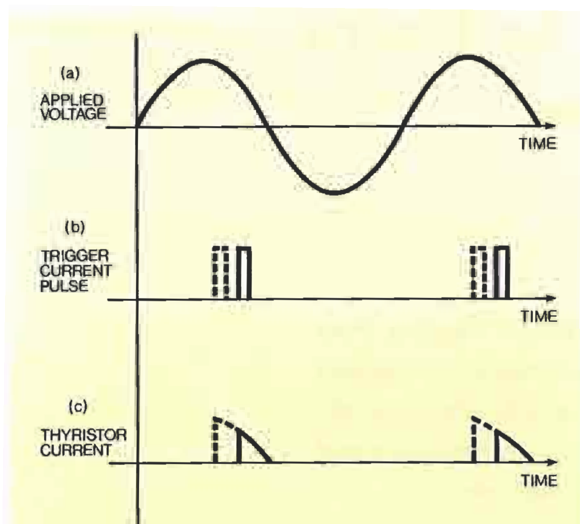

Fig. 13 Power control application of a thyristor. (a) shows the voltage applied to the device, (b) the current pulses which trigger forward conduction and (c) the resulting controlled output current waveform (all to the same time scale). Note that as the trigger pulse is applied earlier (dotted outline), the mean current (dotted and full outlines) increases

of silicon wafers, 50 mm in diameter, after processing, together with some individual devices separated from the wafers. Figure 12 shows thyristors mounted on heat sinks, complete with wires bonded to the device terminals, before and after encapsulation in plastic.

Applications of Thyristors

Thyristors are used in the power control systems of a wide range of AC machines and circuits (7). The basic mechanism of operation is common to all applications and is illustrated in Figure 13. An alternating voltage is applied between the anode and the cathode of the thyristor, suitably chosen to withstand the full voltage excursion in both forward and reverse directions (Figure 13(a)). Figure 13(b) shows a series of current pulses applied to the gate electrode. These pulses can be derived from the voltage signal itself, but their phase relative to that voltage must be variable. The function of the pulses is to trigger the forward-conducting, or 'on' state of the thyristor. The current waveforms of Figure 13(c) are the resulting current flow through the thyristor, with the dotted trigger pulses of Figure 13(b) causing the dotted current waveforms, thus indicating that the earlier the triggering occurs, the greater the mean current flow.

Variable-phase operation allows thyristors to be applied directly to controlling the power input to systems such as light dimmers or electric cookers, the thermal time constant of which is much larger than one period of the AC supply. For systems that respond more rapidly, such as electric motors, it is necessary to smooth the current waveform using filter circuits, but these are simple to provide. Thyristor control of DC motors is achieved by controlling either the armature current, or the field current, or both. This can be applied to motors of almost any size, ranging from those in domestic appliances, to those for electric railway traction, and even to the very large machines which are used in modern steel rolling mills.

The range of application of thyristors in the engineering and consumer markets is almost limitless, and this potential will be increasingly realized as their performance continues to improve and their costs are reduced.

References

1 B. G. Streetman, 'Solid State Electronic Devices', Prentice Hall Inc., Englewood Cliffs, N. J., 1972
2 A. S. Grove, 'Physics and Technology of Semiconductor Devices', John Wiley & Sons Inc., New York, 1967
3 W. Shockley and W. T. Read, *Phys. Rev.*, 1952, **87**, 835-842
4 R. N. Hall, *Phys. Rev.*, 1952, **87**, 387
5 W. R. Wilcox and T. J. LaChappell, *J. Appl. Phys.*, 1964, **35**, 240-246
6 F. A. Trumbore, *Bell Syst. Tech. J.*, 1960, **39**, 205-233
7 F .E. Gentery, F. W. Gutzwiller, N. Holonyak, Jr. and E. E. von Zastrow, 'Semiconductor Controlled Rectifiers: Principles and Applications of *p-n-p-n* Devices', Prentice Hall Inc., Englewood Cliffs, N. J., 1964

Gold-Coated Glass Retains Its Strength

Although the effects of coatings of gold and other metals on the optical properties of glass are well documented, their influence on its mechanical properties have not been closely studied. Past studies indicate that inorganic coatings usually decrease the strength, and therefore safety, of glass, but no satisfactory explanation for this effect was available until the recent publication of work by H. Ishikawa, N. Shinkai and H. Sakata (*J. Mater. Sci.*, 1980, **15**, (2), 483-490) of Asahi Glass Co., Ltd., in Yokohama.

These authors used pre-flawed samples of commercial float glass, coated with chromium, aluminium, gold or silver, and submitted them to three-point static bending tests. Chromium deposits decreased the strength by about 80 per cent and this was more pronounced as the thickness of the coatings was increased from 25 to 250 nm. The effect is attributed to the brittleness of chromium films, their strong bonding to glass and the presence of high tensile stresses in them. Thus, cracks which develop in the metal penetrate into the glass and induce sample failure. With aluminium deposits, the loss of strength is less drastic and it is suspected that cracks are initiated by impurities — possibly oxide particles — at the metal-glass interface. Gold and silver coatings, on the other hand, do not significantly affect the bending strength of glass because tensile stresses in them are low and the metal-glass bond is weak. Hence, cracks in the films do not propagate into the glass.